

Quantum Programming in Polylogarithmic Time*

Florent Ferrari
ENS Lyon, France

Emmanuel Hainry,
Romain Péchoux,
Mário Silva

Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

March 21, 2026

Abstract

Polylogarithmic time delineates a relevant notion of feasibility on several classical computational models such as Boolean circuits or parallel random access machines. As far as the quantum paradigm is concerned, this notion yields the complexity class FBQPOLYLOG of functions approximable in polylogarithmic time with a quantum random access Turing machine. We introduce a quantum programming language with first-order recursive procedures, which provides the first programming language-based characterization of FBQPOLYLOG . Each program computes a function in FBQPOLYLOG (soundness) and, conversely, each function of this complexity class is computed by a program (completeness). We also provide a compilation strategy from programs to uniform families of quantum circuits of polylogarithmic depth and polynomial size, whose set of computed functions is known as QNC , and recover the well-known separation result $\text{FBQPOLYLOG} \subsetneq \text{QNC}$.

Introduction

Motivation. In order to check that quantum programs can be compiled and executed on a quantum computer, one has to design restrictions and constraints implying that quantum programs do not break the laws of quantum mechanics, for example, no-cloning of data and unitarity of operators. In addition, there is a need to tame their complexity in order to ensure their feasibility, that is, the fact that programs do not use too much space and time resources.

Taking inspiration from the classical world, this kind of issue has led to the definition and study of several quantum polynomial time classes. One of the most striking examples of such classes is BQP , the quantum analog of the class of bounded-error probabilistic polynomial time problems BPP . By Yao's theorem [11], BQP corresponds exactly to what can be computed by uniform families of quantum circuits of polynomial size. This class, as well as its extension to functions, FBQP , have been characterized through various means, including function algebras [7] and first-order programs [5].

A natural question is then to study subpolynomial complexity classes. In the literature, polylogarithmic (polylog) time has been introduced and studied on Quantum Random-Access Turing Machine (QRATM) [4]. As in the classical case, this definition uses random-access machines, as opposed to standard quantum Turing machines, because of the sub-linearity of time: although the machine cannot read its entire input, it can access any input bit or qubit. On quantum models, polylog time corresponds to problems that a QRATM can solve in a polylog number of steps, leading to the definition of the complexity class FBQPOLYLOG [8, 9] of functions computable with bounded-error in quantum polylog time.

A main open problem is to design programming languages characterizing such a polylog class abstracting from the low-level considerations (machines, uniformity conditions, etc.) [10].

*This is an extended abstract of [3] presented at MFCS'25.

(Integers)	i	\triangleq	$x \mid k \mid i \pm k \mid i/2 \mid l $
(Booleans)	b	\triangleq	$i \geq i \mid \dots \mid b \wedge b \mid \dots$
(Qubit lists)	l	\triangleq	$\bar{q} \mid l \ominus [i] \mid l^{\boxplus} \mid l^{\boxminus}$
(Qubits)	q	\triangleq	$l[i]$
(Statements)	S	\triangleq	skip ; $q \ast = U^g(i)$; $S \ S$ if b then $\{S\}$ else $\{S\}$ \mid qcase q of $\{0 \rightarrow S, 1 \rightarrow S\}$ \mid call $\text{proc}(l_1, \dots, l_n)$;
(Procedure decl.)	D	\triangleq	$\varepsilon \mid$ decl $\text{proc}(\bar{q}_1, \dots, \bar{q}_n)\{S\}, D$
(Programs)	P	\triangleq	$D :: S$

Figure 1: Syntax of programs

Contributions. This paper makes a first step towards solving this problem. Towards that end, we introduce a quantum programming language with first-order recursive procedures, named PLP for PolyLog Programs (Figure 1), on which we obtain the following results:

- PLP programs are terminating and reversible.
- PLP is sound for FBQPOLYLOG, i.e., each PLP program computes a function in FBQPOLYLOG. Soundness relies on the use of a bounded recursion scheme for procedures to enforce the required polylog time properties, as illustrated by the binary search example.
- PLP is complete for FBQPOLYLOG, i.e., each function of this complexity class is computed by a PLP program. Completeness is shown by a direct encoding of polylog QRATM in PLP.
- PLP is also sound but not complete for QNC. For soundness, we outline a compilation algorithm that from a PLP program and its input size outputs a quantum circuit of polylog depth and polynomial size, i.e., circuits computing functions in QNC. Completeness does not hold as it is well-known that FBQPOLYLOG is strictly included in QNC.

Related Work. A characterization of BQPOLYLOG based on a function algebra has been provided in [8, 9], where a *fast quantum recursion scheme* is used to ensure that programs terminate in polylogarithmic time. Our work employs a similar bounded recursion scheme, using a simple divide-and-conquer strategy on qubits. This characterization can be seen as simpler and more natural approach since an imperative first-order programming language is more accessible to the typical programmer.

First-Order Quantum Programs

Syntax. The considered language is a quantum programming language with first-order recursive procedures whose syntax is provided in Figure 1. There are four basic types τ for expressions: *Integer* expressions are variables x , constant $k \in \mathbb{N}$, arithmetic operations like $i \pm k$ or $i/2$,¹ as well as the size $|l|$ of a list of qubits l . *Boolean* expressions are defined in a standard way. *Qubit lists* are lists of unique (i.e., non-duplicable) qubit pointers. A qubit list expression l is either a variable \bar{q} , the first (respectively second) half l^{\boxplus} (resp. l^{\boxminus}) of the qubit list l , or a list $l \ominus [i]$ where the i -th element of l has been removed. *Qubit* expressions are of the shape $l[i]$, which denotes the i -th qubit in l . We also define syntactic sugar for pointing to the n -th *last* qubit in a list, by defining for any $n \geq 1$, $\bar{q}[-n] \triangleq \bar{q}[|\bar{q}| - n + 1]$.

A program $P \triangleq D :: S$ is defined in Figure 1 as a list of (possibly recursive) procedure declarations D , followed by a program statement S . Statements include the no-op instruction, (single-qubit) unitary application, sequences, conditional, quantum case, and procedure calls. For sake of universality [2], in a unitary application $q \ast = U^g(i)$, the unitary transformation $U^g(i) \in \mathbb{K}^{2 \times 2}$ can take an

¹The semantics of $/2$ will be defined as the ceiling of the result, hence it preserves the set of integers.

integer i and a function $g \in \mathbb{Z} \rightarrow [0, 2\pi)$ as optional arguments², and we omit them when they are of no use.

The quantum conditional **qcase** q **of** $\{0 \rightarrow S_0, 1 \rightarrow S_1\}$ allows branching by executing statements S_0 and S_1 in superposition according to the state of qubit q . When we want to treat cases on multiple qubits, we simply the nested qcases by writing **qcase** q_1, q_2 **of** $\{00 \rightarrow S_{00}, 01 \rightarrow S_{01}, 10 \rightarrow S_{10}, 11 \rightarrow S_{11}\}$.

Semantics. Let $\mathbb{B} \triangleq \{0, 1\}$ denote the set of Booleans and $\mathcal{L}(\mathbb{N})$ denote the set of lists of natural numbers, $[]$ being the empty list. We interpret basic types as follows:

$$\llbracket \text{Integers} \rrbracket \triangleq \mathbb{Z} \quad \llbracket \text{Booleans} \rrbracket \triangleq \mathbb{B} \quad \llbracket \text{Qubit lists} \rrbracket \triangleq \mathcal{L}(\mathbb{N}) \quad \llbracket \text{Qubits} \rrbracket \triangleq \mathbb{N}$$

Qubits are interpreted as integers (pointers) and qubit lists are interpreted as lists of pointers.

Given a program P , let the *length* of P be a function mapping each qubit variable $\bar{q} \in \text{Var}(P)$ to an integer $\text{len}(\bar{q}) \in \mathbb{N}$. We write len_P as a shorthand for $\sum_{\bar{q} \in \text{Var}(P)} \text{len}(\bar{q})$ and $\text{len}_P^{\leq \bar{q}}$ as a shorthand for $\sum_{\bar{q}' \in \text{Var}(P), \bar{q}' < \bar{q}} \text{len}(\bar{q}')$.

A *configuration* c of program P over len_P qubits is of the shape

$$(S, |\psi\rangle, A, f) \in (\text{Statements} \cup \{\top, \perp\}) \times \mathcal{H}_{2^{\text{len}_P}} \times (\text{Var}(P) \rightarrow \mathcal{P}(\mathbb{N})) \times (\text{Var}(P) \rightarrow \mathcal{L}(\mathbb{N})),$$

where \top and \perp are two special symbols denoting termination and error, respectively, where $|\psi\rangle \in \mathcal{H}_{2^{\text{len}_P}}$ is a quantum state, and where, for each qubit list $\bar{q} \in \text{Var}(P)$, $A(\bar{q})$ is the set of qubit pointers accessible from \bar{q} and $f(\bar{q})$ is the list of qubit pointers assigned to \bar{q} .

Given a program $P \triangleq D :: S$, with $n = \text{len}_P$, let Conf_n be the set of configurations over n qubits. The initial configuration in Conf_n on input state $|\psi\rangle \in \mathcal{H}_{2^n}$ is $c_{\text{init}}(|\psi\rangle) \triangleq (S, |\psi\rangle, \bar{q} \mapsto \{1, \dots, \text{len}(\bar{q})\}, \bar{q} \mapsto [1, \dots, \text{len}(\bar{q})])$. A final configuration can be defined in the same way as $c_{\text{final}}(|\psi\rangle) \triangleq (\top, |\psi\rangle, \bar{q} \mapsto \{1, \dots, \text{len}(\bar{q})\}, \bar{q} \mapsto [1, \dots, \text{len}(\bar{q})])$.

Each unitary transformation U of a unitary application $q \ast = U^g(i)$; comes with a function $\llbracket U \rrbracket$ assigning a unitary matrix $\llbracket U \rrbracket(g)(n) \in \mathbb{K}^{2 \times 2}$ to each integer n and function $g \in \mathbb{Z} \rightarrow [0, 2\pi)$. We restrict ourselves to three kinds of gates: the phase gate Ph , rotation gate R_Y and NOT gate NOT , whose semantics is defined as follows:

$$\llbracket Ph \rrbracket(g)(n) \triangleq \begin{pmatrix} 1 & 0 \\ 0 & e^{ig(n)} \end{pmatrix} \quad \llbracket R_Y \rrbracket(g)(n) \triangleq \begin{pmatrix} \cos(g(n)) & -\sin(g(n)) \\ \sin(g(n)) & \cos(g(n)) \end{pmatrix} \quad \llbracket NOT \rrbracket(\cdot)(\cdot) \triangleq \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

The big-step semantics $\cdot \longrightarrow \cdot$ is defined as a relation in $\bigcup_{n \in \mathbb{N}} \text{Conf}_n \times \text{Conf}_n$. We write $\llbracket P \rrbracket(|\psi\rangle) = |\psi'\rangle$, whenever $c_{\text{init}}(|\psi\rangle) \longrightarrow c_{\text{final}}(|\psi'\rangle)$ holds. If the program P terminates on all inputs (i.e., always reaches a final configuration) then $\llbracket P \rrbracket$ is a total function on quantum states.

When an input state is defined by different qubit lists, we denote them in subscript. For instance, for $x, y \in \{0, 1\}^*$ and $m \in \mathbb{N}$, we have that $|x\rangle_{\bar{q}_1} |y\rangle_{\bar{q}_2}$ indicates state $|x\rangle$ given as input to qubit list \bar{q}_1 , state $|y\rangle$ given as input to qubit list \bar{q}_2 .

Example (Binary search). *Let $x \in 0^*1^*2^*$ be a sorted string and \hat{x} denote the encoding of x as a binary given by $\hat{0} \triangleq 00$, $\hat{1} \triangleq 01$, and $\hat{2} \triangleq 10$. Program **SEARCH** in Figure 2 computes the function $\llbracket \text{SEARCH} \rrbracket(|\hat{x}\rangle_{\bar{q}_1} |0\rangle_{\bar{q}_2}) = |\hat{x}\rangle_{\bar{q}_1} |b\rangle_{\bar{q}_2}$, where $b \in \{0, 1\}$ indicates whether x contains a 1 or not.*

Polylogarithmic time restrictions. We now define some restrictions on the admissible programs to guarantee that they terminate in polylogarithmic time (i.e., each procedure cannot perform more than a polylogarithmic number of recursive calls in the input size) and that their total number of sequential procedure calls (i.e., calls that are not in superposition) is bounded polylogarithmically.

Towards that end, we define a relation between procedures to account for recursion. Given a program P , and two procedures proc and proc' appearing in P , let $\text{proc} \sim_P \text{proc}'$ denote that proc and proc' are mutually recursive.

²In the case of quantum polynomial time, Adleman et al. [1] showed how the choice of amplitudes can affect the expressivity of classes such as BQP, requiring the restriction of *polynomial-time approximable complex amplitudes*. How the set of amplitudes influences the class FBQPOLYLOG remains an open question, as discussed in [8, 9], and so we abstain from the use of the entire set of complex numbers and instead use a field \mathbb{K} which may refer to, for instance, polynomial-time approximable complex amplitudes.

```

SEARCH
1 decl search( $\bar{q}_1, \bar{q}_2$ ) {
2   if  $|\bar{q}_1| > 1$  then {
3     qcase  $\bar{q}_1[|\bar{q}_1|/2, |\bar{q}_1|/2 + 1]$  of {
4       00  $\rightarrow$  call search( $\bar{q}_1^{\boxplus} \ominus [1], \bar{q}_2$ );,
5       01  $\rightarrow \bar{q}_2[1] \text{ *= NOT};$ ,
6       10  $\rightarrow$  call search( $\bar{q}_1^{\boxminus} \ominus [-1], \bar{q}_2$ );,
7       11  $\rightarrow$  skip; }
8   } else { skip; }
9 } ::
10 call search( $\bar{q}_1, \bar{q}_2$ );

```

Figure 2: Binary search as an example of a PLP program.

Definition. A program P is said to be recursively halving, denoted $P \in \text{HALF}$, if for each procedure $\text{proc} \in P$ and for each procedure call $\text{call proc}'(l_1, \dots, l_n); \in S^{\text{proc}}$,

if $\text{proc} \sim_P \text{proc}'$ then there are $1 \leq i \leq n$ and l such that l^{\boxplus} or l^{\boxminus} appears in l_i .

This restriction can be viewed as a recursion scheme, which implies a polylogarithmic time bound on programs in HALF by ensuring that in every (mutually) recursive procedure call at least one of the input qubit lists is cut in half.

Now we impose a further condition on the number of sequential (mutually) recursive procedure calls. For that purpose, we define the *width* of a program P in the following way.

Definition. Given a program P , the width of a procedure $\text{proc} \in P$ is defined by $\text{width}_P(\text{proc}) \triangleq w_P^{\text{proc}}(S^{\text{proc}})$ where $w_P^{\text{proc}}(S)$ is defined inductively on statements as follows. For basic instructions, $w_P^{\text{proc}}(\text{skip};) = w_P^{\text{proc}}(q \text{ *= } U^g(i);) \triangleq 0$. For the composition of statements, $w_P^{\text{proc}}(S_0 S_1) \triangleq w_P^{\text{proc}}(S_0) + w_P^{\text{proc}}(S_1)$. Both classical and quantum branching admit the same width, as we have that $w_P^{\text{proc}}(\text{if } b \text{ then } \{S_1\} \text{ else } \{S_0\})$ and $w_P^{\text{proc}}(\text{qcase } q \text{ of } \{0 \rightarrow S_0, 1 \rightarrow S_1\})$ both correspond to $\max(w_P^{\text{proc}}(S_0), w_P^{\text{proc}}(S_1))$. Finally, for procedure calls,

$$w_P^{\text{proc}}(\text{call proc}'(l_1, \dots, l_n);) \triangleq \begin{cases} 1, & \text{if } \text{proc} \sim_P \text{proc}' \\ 0, & \text{otherwise.} \end{cases}$$

The width of a program $\text{width}(P)$ is defined by $\text{width}(P) \triangleq \max_{\text{proc} \in P} \text{width}_P(\text{proc})$. Let $\text{WIDTH}_{\leq 1}$ be defined by $\text{WIDTH}_{\leq 1} \triangleq \{P \mid \text{width}(P) \leq 1\}$.

Definition. The set PLP of PolyLog Programs is defined by $\text{PLP} \triangleq \text{HALF} \cap \text{WIDTH}_{\leq 1}$.

Example. Program SEARCH in Figure 2 can be shown to be in PLP . This program only contains one procedure search , that is recursive (as $\text{search} \sim_{\text{SEARCH}} \text{search}$ holds). It is then easy to check that $\text{SEARCH} \in \text{HALF}$ as the recursive procedure calls are performed either on qubit lists \bar{q}^{\boxplus} or \bar{q}^{\boxminus} . Furthermore, we verify that $\text{SEARCH} \in \text{WIDTH}_{\leq 1}$ by computing $\text{width}_{\text{SEARCH}}(\text{search})$:

$$\begin{aligned} w_{\text{SEARCH}}^{\text{search}}(S^{\text{search}}) &= \max(w_{\text{SEARCH}}^{\text{search}}(\text{qcase } \dots), 0) = \max(w_{\text{SEARCH}}^{\text{search}}(\text{call search}(\bar{q}_1^{\boxplus} \ominus [1], \bar{q}_2);), \\ &\quad w_{\text{SEARCH}}^{\text{search}}(\text{call search}(\bar{q}_1^{\boxminus} \ominus [-1], \bar{q}_2);)) \\ &= \max(1, 1) = 1. \end{aligned}$$

A Characterization of Quantum Polylog Time

PLP characterizes exactly the functions in FBQPOLYLOG , the class of quantum polylog time approximable functions. That is, programs in PLP compute functions in FBQPOLYLOG (Soundness) and, reciprocally, for any function in FBQPOLYLOG , there exists a PLP program that computes it (Completeness).

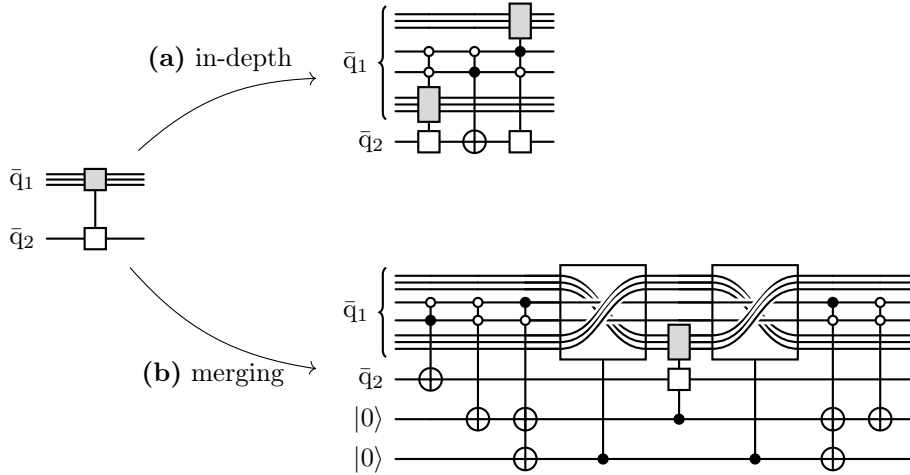


Figure 3: Compilation strategies for `search` defined in Figure 2.

Main result. We denote by $\llbracket \text{PLP} \rrbracket$ the set of functions computed by programs in PLP. That is $\llbracket \text{PLP} \rrbracket \triangleq \{ \llbracket P \rrbracket \mid P \in \text{PLP} \}$. A program P approximates function $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ with probability $p \in [0, 1]$ if $\forall x \in \{0, 1\}^*, |\langle f(x) | \llbracket P \rrbracket(x) \rangle|^2 \geq p$, in other words if for all input, the output of P coincides with f with probability at least p . The set of functions that can be approximated with probability at least p is denoted by $\llbracket \text{PLP} \rrbracket_{\geq p}$.

Theorem (Soundness & Completeness). $\llbracket \text{PLP} \rrbracket_{\geq \frac{2}{3}} = \text{FBQPOLYLOG}$.

Proof. Soundness is proved via the fact that the `HALF` and `WIDTH≥1` restrictions imply a polylogarithmic bound for the depth of the call tree and a logarithmic bound on the degree of this tree. Then we show that if some PLP program P approximates f , then the poly-logarithmic time QRATM simulating P also approximates f and guarantees that $f \in \text{FBQPOLYLOG}$. Conversely, for completeness, given a polylog time QRATM M , we define a PLP program that simulates M . \square

Circuit compilation. The compilation strategy takes inspiration from [5, 6] which in particular uses ancillas to factorize the circuits representing procedure calls in branches so as to prevent exponential blow-up of the circuit size. Their technique is called *anchoring and merging* as when a procedure call is first encountered, an ancilla is associated to this call (*anchoring*), and when a subsequent call to this procedure happens with an input of the same size, this second call is then merged with the first (*merging*). This way, instead of doubling the size of the circuit whenever recursive calls appear in separate branches of a `qcase`, as in program `SEARCH` of Figure 2, the size grows linearly in the number of nested recursive calls, hence preventing the exponential blow-up in complexity from the use of the quantum control statement [12].

Figure 3 exemplifies this phenomenon on the `SEARCH` program: the circuit on the left represented by a grey and white box the circuit for the `search` procedure applied to \bar{q}_1, \bar{q}_2 . Since this procedure has two calls to itself, its natural compilation gives an in-depth duplication of the calls. The anchoring/merging process entails a single recursive compilation at the price of an overhead in terms of ancillary qubits and permutations. Importantly, controlled permutations can be performed in logarithmic depth, which ensures that the final circuit size and depth bounds place all PLP programs in QNC.

Theorem (Compilation). *Given a PLP program P , and input size $n = \sum_{\bar{q} \in \text{Var}(P)} |\bar{q}|$, the quantum circuit produced by the compilation process is of size $O(n \text{polylog}(n))$ and depth $O(\text{polylog}(n))$.*

References

- [1] Leonard M. Adleman, Jonathan DeMarras, and Ming-Deh A. Huang. Quantum computability. *SIAM Journal on Computing*, 26(5):1524–1540, 1997.
- [2] P. Oscar Boykin, Tal Mor, Matthew Pulver, Vwani Roychowdhury, and Farrokh Vatan. On universal and fault-tolerant quantum computing, 1999.
- [3] Florent Ferrari, Emmanuel Hainry, Romain Péchoux, and Mário Silva. Quantum Programming in Polylogarithmic Time. In Paweł Gawrychowski, Filip Mazowiecki, and Michał Skrzypczak, editors, *50th International Symposium on Mathematical Foundations of Computer Science (MFCS 2025)*, volume 345 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 47:1–47:17, Dagstuhl, Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [4] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Quantum random access memory. *Phys. Rev. Lett.*, 100:160501, Apr 2008.
- [5] Emmanuel Hainry, Romain Péchoux, and Mário Silva. A programming language characterizing quantum polynomial time. In *Foundations of Software Science and Computation Structures*, pages 156–175. Springer, 2023.
- [6] Emmanuel Hainry, Romain Péchoux, and Mário Silva. Branch Sequentialization in Quantum Polytime. In Maribel Fernández, editor, *10th International Conference on Formal Structures for Computation and Deduction (FSCD 2025)*, volume 337 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 22:1–22:22, Dagstuhl, Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [7] Tomoyuki Yamakami. A schematic definition of quantum polynomial time computability. *J. Symb. Log.*, 85(4):1546–1587, 2020.
- [8] Tomoyuki Yamakami. Expressing power of elementary quantum recursion schemes for quantum logarithmic-time computability. In *Logic, Language, Information, and Computation (WoLLIC 2022)*, pages 88–104. Springer, 2022.
- [9] Tomoyuki Yamakami. Elementary quantum recursion schemes that capture quantum polylogarithmic-time computability of quantum functions. *Mathematical Structures in Computer Science*, 34(7):710–745, August 2024.
- [10] Tomoyuki Yamakami. Quantum first-order logics that capture logarithmic-time/space quantum computability. In Ludovic Levy Patey, Elaine Pimentel, Lorenzo Galeotti, and Florin Manea, editors, *CiE 2024*, pages 311–323. Springer, 2024.
- [11] Andrew Chi-Chih Yao. Quantum circuit complexity. In *Proceedings of 1993 IEEE 34th Annual Foundations of Computer Science*, pages 352–361, 1993.
- [12] Charles Yuan and Michael Carbin. Tower: Data structures in quantum superposition. *Proceedings of the ACM on Programming Languages*, 6(OOPSLA2):259–288, Oct 2022.